

UNITED STATES PATENT APPLICATION

FOR

MULTI-SESSION NO QUERY RESTORE

BY

ROB EDWARDS, BARRY FRUCHMAN, AND OMAR VARGAS

10094797-030600

MULTI-SESSION NO QUERY RESTORE

FIELD OF THE INVENTION

5 This invention relates to data restoration, and in particular to coordinated data restoration in a computer network system in which a plurality of client systems may participate in the same restore.

BACKGROUND OF THE INVENTION

10 Data is the underlying resource on which all computing processes are based. The volume of this data is ever increasing in various business applications, and each bit of data may be critical to business functions. In dealing with such data, many businesses arrange their personal computers (PCs) in a client/server network. The PCs and any associated programs typically function as the client in that it requests files or services. A server, which may be anything from another PC to a mainframe, handles the request from the client and supplies the requested data or service to the client. Such client/server networks may communicate with each other over an associated local area network (LAN).

15 To address the volume and importance of storing such data on client/server networks, storage area networks (SANs) have emerged to free up bandwidth on such LANs and to provide
20 storage and related storage services to clients of one or more client systems such as backup and restoration functions. A SAN is a dedicated network separate from LANs and wide area networks (WANs) which interconnects storage devices to one or more servers and to a plurality of clients and/or client systems in a related network.

Storage devices are a place to keep and retrieve data on a long-term basis. Each storage
25 device includes some storage medium which physically stores the data such magnetic tape,
optical disks, hard disks, and floppy disks. Storage media can also be arranged in a variety of
ways including a redundant array of independent or inexpensive disks (RAIDs) which typically
function as one of the storage devices in a SAN.

SANs also often have high interconnect data rates (gigabits/second) between member
30 storage devices and are highly scalable. SANs can be interconnected with similar elements as in
LANs and WANs, e.g., routers, hubs, switches, and gateways. A SAN may be local or extend
over geographic distances.

A storage management server or servers may also be utilized to control the storage
devices and keep track of the data that the plurality of clients have stored on the plurality of
35 storage devices coupled to a common SAN. The storage management server may also be
utilized in data restoration efforts. Data restoration permits clients to copy a version of a backup
file or files stored on any one of the plurality of storage devices.

However, such data restoration efforts typically require manual restoration and
partitioning and do not permit a plurality of client systems to automatically participate in the
40 same restore effort. Accordingly, there is a need in the art for a system and method for
optimizing data restoration in a SAN environment where multiple client systems may be
involved in a single restore which allows for coordinated access to an associated storage pool of
data.

BRIEF SUMMARY OF THE INVENTION

A method of restoring data in a computer network system wherein a plurality of client systems have access to a storage pool coupled to an associated storage area network (SAN)

50 consistent with the invention includes: requesting a restore wherein each of the plurality of client systems may participate in the restore; and coordinating access to the data stored in the storage pool by tracking a plurality of data portions of the data to be restored and by blocking access to each of the plurality of data portions that have been restored by one of the plurality of client systems to avoid duplicative restoration efforts.

55 A computer network system for restoring data consistent with the invention includes: a plurality of client systems; a storage pool coupled to the plurality of client systems through a storage area network (SAN); and a storage management server coupled to the plurality of client systems through the SAN, wherein the storage management server is configured to coordinate access to the data stored in the storage pool by tracking a plurality of data portions of the data to be restored and by blocking access to each of the plurality of data portions that have been
60 restored by one of the plurality of client systems to avoid duplicative restoration efforts.

BRIEF DESCRIPTION OF THE DRAWINGS

65 For a better understanding of the present invention, together with other objects, features and advantages, reference should be made to the following detailed description which should be read in conjunction with the following figures wherein like numerals represent like parts:

FIG. 1 is a block diagram of a computer network system including a plurality of client systems with access to a storage pool of a SAN; and

FIG. 2 is an exemplary master restore table consistent with the present invention.

DETAILED DESCRIPTION

Turning to FIG. 1, a block diagram of a computer network system 100 including a storage management server 102 configured to create and store a master restore table consistent with the present invention is illustrated. The computer network system 100 may contain a plurality of client systems 103, 115. Although only two client systems 103, 115 are illustrated for clarity, those skilled in the art will recognize that a computer network system 100 consistent with the invention may contain any number of client systems.

Each client system may contain a plurality of clients. For instance, a first client system 103 may contain its associated plurality of clients 107, 109, 111 and a second client system 115 may similarly contain its associated plurality of clients 117, 119, 121. Each client may be a PC or software on such PCs. Clients in each client system 103, 115 are coupled directly to each other in any number of fashions known to those skilled in the art, but are not coupled directly to other clients in other client systems. In addition, those skilled in the art will recognize any number of clients may be present in a computer network system 100 consistent with the invention.

One or more of the clients 107, 109, 111 or 117, 119, 121 may be equipped with storage software, e.g., a storage agent, enabling the client with such software to communicate data to be stored in associated storage pool 134 directly over the SAN 106. The storage pool 134 includes a plurality of storage devices 112, 114, 116. Any number of such storage devices 112, 114, 116 may be present in a computer network system 100 consistent with the invention. Each storage device 112, 114, 116 includes some storage medium which physically stores the data such magnetic tape, optical disks, hard disks, floppy disks, or the like.

Such data transfers between client systems 103, 115 and the storage pool 134 normally occur without moving data through the LAN 104, the LAN 108, or the storage management server 102. Accordingly, LAN communication bandwidth can be freed up for a variety of uses and a corresponding decreased load on the storage management server 102 allows it to support a greater number of simultaneous client connections.

The client systems 103, 115 may continue to use the LAN 104 or 108 connections to the storage management server 102 to exchange control information over the LAN 104 or 108 to the server 102, such as policy information and data about the objects that are backed up. If a failure occurs and data cannot be transferred from the client systems 103, 115 over the SAN 106 to the storage pool 134, then the client systems 103, 115 may use an alternate path via the LAN 104 or LAN 108 to make such a connection.

Advantageously, the computer network system 100 may have a storage management server 102 configured to construct a master restore table consistent with the invention to optimize data restore efforts among a plurality of client systems 103, 115. As data from each client of each client system 103, 115 is backed up and stored in one or more storage devices 112, 114, 116 of the storage pool 134, the storage management server 102 receives control information indicating the stored location of object data.

When a client having appropriate storage software from a client system 103, 115 requests a restore, the server 102 constructs a master restore table associated with that particular restore request. In general, the master restore table includes portions of data to be restored and an associated location of those portions of data in the storage pool 134. As the master restore table is being constructed, data restoration efforts may begin without waiting for the table to be completed.

Advantageously, as described more fully below with reference to FIG. 2, clients 107, 109, 111 and 117, 119, 121 of a plurality of client systems 103, 115 may participate in a restore effort by utilizing a master restore table to coordinate access to data stored in the storage pool 134 such that duplicative restoration efforts from a plurality of client systems 103, 115 can be avoided.

Turning to FIG. 2, an exemplary master restore table 200 consistent with the invention that may be created and temporarily stored in a storage management server 102 or any device common to the system 100 is illustrated. The table 200 generally is used to track portions of data to be restored and the associated restore media from the various storage devices 112, 114, 116 of the storage pool 134 where such portions of data are located. A host of clients 107, 109, 111 and 117, 119, 121 and client systems 103, 115 may be able to access the table 200 to optimize restore efforts from a plurality of client systems 103, 115. Although described in terms of columns and rows, a master restore table consistent with the invention may take a variety of forms.

The master restore table 200 may include a plurality of columns including: a first column 202 detailing the portions of data to be restored; a second column 204 detailing the location of such portions of data on associated media from various storage devices 112, 114, 116 of the storage pool 134; a third column 206 detailing a LAN-free path for accessing the associated storage media if such path exists, a fourth column 208 detailing a server-free path for accessing the associated storage media if such path exists, and a fifth column 210 detailing the status of whether an associated storage media has been processed for a given portion of data.

Advantageously then, the fifth column tracks the portions of data that have been processed by any one client in any one plurality of client systems 103, 115. In this way, clients or restore processes from any client system 103, 115 would be blocked from restoring a portion of that data

that had already been processed or restored. Accordingly, duplicative restoration efforts are automatically avoided.

In addition, the portions of data may be automatically partitioned into a plurality of sub-tables, e.g., a first-sub-table 215, a second sub-table 217, and a third sub-table 219 are illustrated although any number of sub-tables may be necessary depending on the volume and location of the objects to be restored. Advantageously, the partitioning of the portions of data into the various sub-tables 215, 217, 219 occurs automatically, and the partitioning may be based on predetermined criteria such as the location of the portions of data on associated storage devices 112, 114, 116 in the storage pool 134. In addition, restoration of objects in the first sub-table 215 may occur as future sub-tables are still being constructed. The table 200 also permits a restore to continue after it has been interrupted since the fifth column 210 keeps track of which objects have been processed for the restore.

The master restoration table also coordinates restoration of a plurality of data portions concurrently from a plurality of storage devices 112, 114, 116 in order to efficiently restore lost data. Data to be restored may be provided to a target restoration device, which may be any device, e.g., a common server, which is accessible to the plurality of client systems 103, 114.

Restoration may also take place during one or a plurality of sessions. A session is a period of time in which a user or client can communicate with an associated server, e.g., the storage management server 102, to perform a backup, archive, or restore request. In this way, restoration software in one or more concurrent sessions from one or more client systems 103, 115 can be run in order to optimize a restoration effort.

Clients 107, 109, 111 or 117, 119, 121 with appropriate storage software may gain access to the restore table 200 by means of a restore token, e.g., a predetermined sequence of bits, which

uniquely identifies the table that the server has built for a particular restore effort. In addition, each client 107, 109, 111 or 117, 119, 121 may gain access the restore table multiple times simultaneously. When the restore effort for a particular restore request is complete, the original
165 initiating client communicates with the storage management server 102 to delete the master restore table. The token may then be available for other clients to use.

The embodiments that have been described herein, however, are but some of the several which utilize this invention and are set forth here by way of illustration but not of limitation. It is obvious that many other embodiments, which will be readily apparent to those skilled in the art,
170 may be made without departing materially from the spirit and scope of the invention.